arXiv:1406.1158v2 [cs.DS] 11 Jan 2015

Kernelization lower bound for Permutation Pattern Matching*

Ivan Bliznets†, Marek Cygan‡, Paweł Komosa§, Lukáš Mach¶

January 13, 2015

Abstract

A permutation π contains a permutation σ as a pattern if it contains a subsequence of length $|\sigma|$ whose elements are in the same relative order as in the permutation σ. This notion plays a major role in enumerative combinatorics. We prove that the problem does not have a polynomial kernel (under the widely believed complexity assumption $\mathsf{NP} \not\subseteq \text{co-}\mathsf{NP}/\text{poly}$) by introducing a new polynomial reduction from the clique problem to permutation pattern matching.

1 Introduction

Counting permutations of size n avoiding a fixed pattern is an established and active area of enumerative combinatorics. Knuth [12] has shown that the number of permutations avoiding $(2, 3, 1)$ is the n^{th} Catalan number. Various choices of prohibited patterns have been studied among others by Lovász [13], Rotem [15], and Simion and Schmid [17]. This culminated in the Stanley-Wilf conjecture stating that for every fixed prohibited pattern, the number of permutations of length n avoiding it can be bounded by c^n for some constant c. Klazar [11] reduced the question to the Füredi-Hajnal conjecture, which was ultimately proved by Marcus and Tardos in 2004 [14].

Wilf [18] also posed the algorithmic question of whether detecting a given pattern (of length ℓ) in a given permutation (of length n) can be done in subexponential time. Subsequently, the problem was shown to be NP-hard in [6]. Ahal and Rabinovich have obtained an $\mathcal{O}(n^{0.47\ell+o(\ell)})$ time algorithm [1]. Fast algorithms have been found for certain restricted versions of the problem [10, 16].

*Partially supported by Warsaw Center of Mathematics and Computer Science.

†St. Petersburg Department of Steklov Institute of Mathematics. E-mail: iabliznets@gmail.com. Partially supported by the Government of the Russian Federation (grant 14.Z50.31.0030).

‡Institute of Informatics, University of Warsaw, Poland. E-mail: cygan@mimuw.edu.pl. Partially supported by NCN grant DEC-2012/05/D/ST6/03214.

§Institute of Informatics, University of Warsaw, Poland. E-mail: kompaw01@gmail.com.

¶DIMAP and Department of Computer Science, University of Warwick, United Kingdom. E-mail: lukas.mach@gmail.com. This author has also recieved funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 259385.

Pattern matching has also received interest in the context of parameterized complexity. Several groups of researchers have obtained W[1]-hardness results for generalizations of the problem [4, 9]. In [5] it was shown that the problem is in FPT when parameterized by the number of runs (maximal monotonic consecutive subsequences) in the target permutation. The authors of [5] raise the issue of whether their problem has a polynomial size kernel as an open problem. The central question of whether the problem is in FPT when parameterized by ℓ has been resolved by Guillemot and Marx [9], who obtained an algorithm with runtime of $2^{\mathcal{O}(\ell^2 \log \ell)} \cdot n$. This implies the existence of a kernel for the problem. Obtaining kernel size lower bounds was posed as an open question during a plenary talk at *Permutation Patterns 2013* by Stéphane Vialette.

We prove that the permutation pattern problem under the standard parameterization by ℓ does not have a polynomial size kernel, assuming NP $\not\subseteq$ co-NP/poly.

2 Preliminaries

The set $\{i, i+1, \ldots, j-1, j\}$ is denoted by $[i, j]$. We set $[n] := [1, n]$. A permutation π is a bijection from $[n]$ to $[n]$. The value $\pi(i)$ is called *the entry of π at position i*. We use $|\pi|$ to denote the size of the domain of π. Two common representations of a permutation π are used: the vector $(\pi(1), \pi(2), \ldots, \pi(n))$ and the corresponding permutation matrix. The latter is a $|\pi| \times |\pi|$ binary matrix with 1-entries precisely on coordinates $(\pi(i), i)$. A vector obtained from the vector representation by omitting some entries is *a subsequence of the permutation*. Such subsequence is *a consecutive subsequence* if it contains precisely the entries with indexes from $[i, j]$ for some $i, j \in \mathbf{N}$. We use $\pi[i, j]$ to denote the set of entries $\{\pi(i), \pi(i+1), \ldots, \pi(j)\}$. *A monotonic subsequence* is a subsequence whose entries form a monotonic sequence. *A run* is a maximal monotonic consecutive subsequence. For example, $(4, 5, 3, 1, 2)$ contains a (decreasing) run of length 3.

The key notion of a permutation pattern is introduced below:

Definition 1. *A permutation σ on the set $[l]$ is* a pattern of a permutation π *on the set $[n]$ if there exists an increasing function $\phi : [l] \to [n]$ such that*

$$\forall x, y \in [l] : \sigma(x) < \sigma(y) \text{ if and only if } \pi(\phi(x)) < \pi(\phi(y)).$$

We say that the function ϕ certifies the pattern.

A parameterized problem is a language $Q \subseteq \Sigma^* \times \mathbf{N}$, where Σ is a fixed alphabet. The value k of the instance $(x, k) \in Q$ is its *parameter*. Permutation Pattern Matching is the following parameterized algorithmic problem:

Input: a permutation σ on $[\ell]$, a permutation π on $[n]$.
Parameter: ℓ.
Question: is σ a pattern of π?

A parameterized problem Q is in FPT if there is an algorithm deciding $(x, k) \in Q$ in time $f(k)|x|^{\mathcal{O}(1)}$, where f is a computable function.

Definition 2. A kernelization algorithm for a parameterized problem Q *is an algorithm that given an instance* $(x, k) \in \Sigma^* \times \mathbf{N}$ *produces in* $p(|x| + k)$ *steps an instance* (x', k') *such that*

1. $(x, k) \in Q \Leftrightarrow (x', k') \in Q$ *and*
2. $|x'|, k' \leq f(k)$,

where $p(\cdot)$ *is a polynomial and* $f(\cdot)$ *a computable function.*

If there is a kernelization algorithm for Q, we say that Q *has a kernel*. If the function $f(\cdot)$ in the above definition can be bounded by a polynomial, we say that Q *has a polynomial kernel*.

We utilize the standard machinery of Bodlaender et al. [3], which builds on [2, 7], to derive kernelization lower-bounds under the widely believed complexity assumption $\mathsf{NP} \not\subseteq \text{co-}\mathsf{NP}/\text{poly}$. A failure of this assumption would imply that the polynomial hierarchy collapses to the third level. Some basic definitions are necessary:

Definition 3 (Bodlaender et al. [3])**.** *An equivalence relation* R *on* Σ^* *is called* a polynomial equivalence relation *if the following two conditions hold:*

1. *There is an algorithm that given two strings* $x, y \in \Sigma^*$ *decides whether* x *and* y *belong to the same equivalence class in* $(|x| + |y|)^{\mathcal{O}(1)}$ *time.*
2. *For any finite set* $S \in \Sigma^*$ *the equivalence relation* R *partitions the elements of* S *into at most* $(\max_{x \in S} |x|)^{\mathcal{O}(1)}$ *equivalence classes.*

An example of such a relation is the grouping of instances of the same size.

Definition 4 (Bodlaender et al. [3])**.** *Let* $L \subseteq \Sigma^*$ *be a language and let* $Q \subseteq \Sigma^* \times \mathbf{N}$ *be a parameterized problem. We say that* L cross-composes into Q *if there is a polynomial equivalence relation* R *and an algorithm which, given* t *strings* $x_1, x_2, \ldots, x_t$ *belonging to the same equivalence class of* R*, computes an instance* $(x^*, k^*) \in \Sigma^* \times \mathbf{N}$ *in time polynomial in* $\sum_{i=1}^{t} |x_i|$ *such that:*

1. $(x^*, k^*) \in Q \Leftrightarrow x_i \in L$ *for some* $1 \leq i \leq t$,
2. k^* *is bounded by a polynomial in* $\max_{i=1}^{t} |x_i| + \log t$.

Theorem 5 (Bodlaender et al. [3])**.** *Let* $L \subseteq \Sigma^*$ *be an* NP*-hard language. If* L *cross-composes into a parameterized problem* Q *and* Q *has a polynomial kernel then* $\mathsf{NP} \subseteq$ *co-*NP*/poly.*

3 Kernelization lower bound for Permutation Pattern Matching

We show that the Permutation Pattern Matching problem is unlikely to have a polynomial kernel:

Theorem 6. *Unless NP $\subseteq$ co-NP/poly, the* Permutation Pattern Matching *problem does not have a polynomial kernel.*

We prove Theorem 6 using Theorem 5. However, this requires a polynomial-time reduction that allows cross-composition without significantly increasing the parameter value. Reductions described in the literature [4,6] have resisted our attempts to apply the framework. Therefore, we introduce a new NP-hardness proof that directly leads to a cross-composition. The new reduction is from the well known Clique problem.

Let us first define encoding $\pi_z(G)$ taking a graph G and $z \in \mathbf{N}$ and producing a permutation. The key property of the encoding is that for any G, H we have $G \subseteq H$ if and only if $\pi_z(G)$ is a pattern of $\pi_z(H)$ for a particular choice of z. This allows us to express the Clique problem in terms of Permutation Pattern Matching. The encoding permutation itself consists of two types of entries: encoding entries and separating entries. The former ones encode the edges of G. The latter form decreasing runs used to separate encoding entries corresponding to different vertices. The mapping $\pi_z(\cdot)$ can be seen as an embedding of the upper-triangular submatrix of the adjacency matrix of the input graph into a permutation. The separating runs mark where each row and column begins and ends; the encoding entries determine where the 1-entries of the matrix are. This mapping, along with some notation introduced below, is illustrated in Figure 1.

We start constructing $\pi_z(G)$ by imposing a total order on $V(G)$ placing vertices from the same connected component of G consecutively. Thus, we can assume that $V(G) = [n]$ and set

$$\begin{aligned} N_G^+(v) &:= \{u : u > v \wedge \{u, v\} \in E(G)\}, \\ N_G^-(v) &:= \{u : u < v \wedge \{u, v\} \in E(G)\}, \\ \deg_G^+(v) &:= |N_G^+(v)|, \\ \deg_G^-(v) &:= |N_G^-(v)|. \end{aligned}$$

We call the vertices from $N_G^+(v)$ and $N_G^-(v)$ the right-neighbours and left-neighbours of v, respectively. The overall structure of the permutation $\pi_z(G)$ is as follows. It starts with a decreasing run of length z, continues with the entries encoding the vertex 1 (i.e., encoding $N_G^+(1)$), which is then followed by another decreasing run of length z. This finishes the part of the permutation dedicated to the vertex 1 and the segment for the vertex 2 begins. Again, it starts with another decreasing run of length z, continues with the encoding entries of $N_G^+(2)$, and is finished by a decreasing run of length z. This continues for all vertices of G. Note that for each vertex v there is a pair of decreasing runs immediately surrounding the entries encoding $N_G^+(v)$, one from left and one from right. These are called the left and right separating runs of v, respectively.

To facilitate the formal definition of $\pi_z(G)$, we begin by introducing a notation for important positions and values of the resulting permutation's entries. This includes the positions where the abovementioned runs start, the values with which they start, or the positions where the parts encoding $N_G^+(v)$, for individual choices of v, start.

We use $p_L(v)$ and $p_R(v)$ as a shorthand for the positions on which the left and right separating run of v starts, respectively. The first position of the segment encoding

$N_G^+(v)$ is denoted by $p_M(v)$. This is illustrated in Figure 1. Specifically, we set $p_L(1) := 1, p_M(1) := z+1$, and $p_R(1) := z+1+\deg_G^+(1)$. For $v \geq 2$, we have:

$$\begin{aligned} p_L(v) &:= p_R(v-1) + z, \\ p_M(v) &:= p_L(v) + z, \\ p_R(v) &:= p_M(v) + \deg_G^+(v). \end{aligned}$$

We also introduce notation for the *values* used by the separating runs. The left separating run of v starts at the position $p_L(v)$ with the value $q_L(v)$. The right separating run starts at $p_R(v)$ with the value $q_R(v)$. Finally, $q_M(v)$ is the least value used by the encoding entries of vertices from $N_G^-(v)$ to determine their connection to v. (Specifically, vertices of $N_G^-(v)$ use the values $[q_M(v), q_M(v)+\deg_G^-(v)-1]$ to encode this. If $\deg_G^-(v)$ is zero, the value $q_M(v)$ is actually not used.) We set $q_L(1) := 2z, q_M(1) := z+1$, and

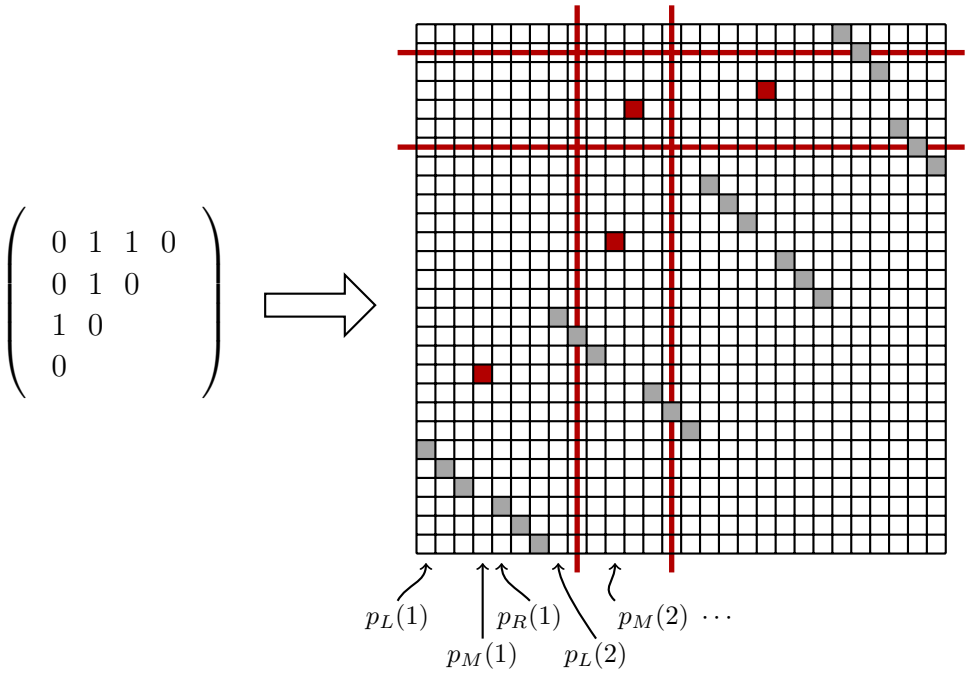

Figure 1: Left part shows the upper triangular submatrix of the adjacency matrix of a graph $G = \big(\{1,2,3,4\}, \big\{\{1,2\},\{2,3\},\{2,4\},\{3,4\}\big\}\big)$. The right part shows the permutation matrix representation of its encoding permutation $\pi_3(G)$. The columns of both matrices are indexed from left to right, the rows from bottom to top. Thus, the $(1,1)$ entry of the matrix is in the bottom-left corner. White positions of the grid on the right correspond to 0-entries of the permutation matrix, non-white positions are 1-entries. Separating runs are colored in light gray, encoding entries in dark red. Note the one-to-one correspondence between the 1-entries of the matrix on the left with the encoding entries of the permutation matrix. Horizontal red lines represent *values* attained at positions $L(4)$ and $R(4)$. Vertical red lines denote *indexes* $L(2)$ and $R(2)$. Note that these four red lines induce a rectangle with a single 1-entry. This encodes the 1-entry in the top-most position of the second column of the adjacency matrix. Arrows below the permutation matrix illustrate the notation $p_L(\cdot), p_M(\cdot)$, and $p_R(\cdot)$.

$q_R(1) := z$. For $v \geq 2$, let

$$\begin{aligned} q_R(v) &:= q_L(v-1) + z, \\ q_M(v) &:= q_R(v) + 1, \\ q_L(v) &:= q_M(v) + z + \deg_G^-(v) - 1. \end{aligned}$$

We now define the values of $\pi = \pi_z(G)$. For each v, we introduce a decreasing run of length z starting at the position $p_L(v)$:

$$\begin{aligned} \pi(p_L(v)) &:= q_L(v), \\ \pi(p_L(v)+1) &:= q_L(v) - 1, \\ \pi(p_L(v)+2) &:= q_L(v) - 2, \\ &\dots \\ \pi(p_L(v)+z-1) &:= q_L(v) - (z-1). \end{aligned}$$

We also insert a decreasing run which starts at the position $p_R(v)$ with the value $q_R(v)$:

$$\begin{aligned} \pi(p_R(v)) &:= q_R(v), \\ \pi(p_R(v)+1) &:= q_R(v) - 1, \\ \pi(p_R(v)+2) &:= q_R(v) - 2, \\ &\dots \\ \pi(p_R(v)+z-1) &:= q_R(v) - (z-1). \end{aligned}$$

This establishes the entries represented by gray squares in Figure 1.

The remaining values are used to encode the edges of G. The neighbourhood $N_G^+(v)$ is encoded by an increasing run on positions $p_M(v), p_M(v)+1, \dots, p_M(v)+|N_G^+(v)|-1$. We fix a vertex $v \in V(G)$ and iterate through the neighbours $\{u_1, u_2, \dots, u_k\} = N_G^+(v)$. Assume $u_1 < u_2 < \dots < u_k$. For $i \in [k]$, we set:

$$\pi(p_M(v) + i - 1) := q_M(u_i) + \ell(v, u_i), \tag{1}$$

where $\ell(v, u_i) = |\{w : w < v \wedge \{w, u_i\} \in E(G)\}|$. The term $\ell(v, u_i)$ ensures that no value in π is repeated.

The above procedure is carried out for each $v \in V(G)$. This finishes the construction of $\pi_z(G)$. We now provide two observations.

Observation 7. *For any graph G and $z \in \mathbf{N}$ the function $\pi_z(G)$ is a permutation.*

Proof. Let $\pi := \pi_z(G)$. It is straightforward to verify that π is a mapping from $[p]$ to $[p]$, for $p = 2zn + |E(G)|$. It remains to show that π is injective, i.e. that there is no pair of distinct indexes i, j such that $\pi(i) = \pi(j)$. It is easily seen that such i and j cannot both be an index of an entry forming a separating run, since the separating runs are explicitly constructed so that the sets of their values are disjoint. For each vertex v there are exactly $\deg_G^-(v)$ values between the values of its left and right separating run. These values are used to encode the $\deg_G^-(v)$ edges connecting v to its left-neighbours. The left-most neighbour is using the least value, the subsequent vertices are using values that increase by 1 with each neighbour (cf. the term $\ell(v, u_i)$

in equation (1)). Therefore, we have neither a collision between a separating entry and an encoding entry nor a collision between two entries encoding $N_G^-(v)$ for the same v. Finally, it can be easily seen that the sets of values encoding $N_G^-(v)$ are pairwise disjoint for different choices of v. □

Observation 8. *For any choice of $z \in \mathbf{N}$ and any choice of $u, v \in V(G)$, there is at most one 1-entry of $\pi_z(G)$ with an index in $[p_M(u), p_R(u) - 1]$ and value in $[q_M(v), q_M(v) + \deg_G^-(v) - 1]$. Furthermore, there is an edge between vertices $u, v \in V(G), u < v$ if and only if there is exactly one such 1-entry.*

Proof. The entries of $\pi := \pi_z(G)$ with indexes from $[p_M(u), p_R(u) - 1]$ encode the neighbourhood of the vertex u, i.e., $N_G^+(u)$. For each neighbour v from $N_G^+(u)$, we insert a single entry with value from $[q_M(v), q_M(v) + \deg_G^-(v) - 1]$. This is because in equation (1) the term $\ell(v, u_i)$ is always strictly less than $\deg_G^-(v)$. This implies both parts of the observation. □

For the purpose of the proof of the lemma below, we define the following:

$$\begin{aligned} C(v) &:= [p_M(v), p_R(v) - 1], \\ L(v) &:= p_L(v) + \lfloor \tfrac{z}{2} \rfloor, \\ R(v) &:= p_R(v) + \lfloor \tfrac{z}{2} \rfloor. \end{aligned}$$

Therefore, $C(v)$ is the set of entries of π encoding the vertex v, $L(v)$ denotes the middle entry of the left separating run of v, and $R(v)$ is the middle entry of the right separating run of v. Once more, Figure 1 illustrates the notation.

The following lemma implies NP-hardness of Permutation Pattern Matching:

Lemma 9. *For every graph G without isolated vertices, K_l is a subgraph of G if and only if $\pi_z(K_l)$ is a pattern of $\pi_z(G)$, for $z = 4n' + 4$, where n' is the number of vertices in the largest connected component of G.*

Proof. We let $\sigma := \pi_z(K_l)$ and $\pi := \pi_z(G)$.

If G contains a clique K_l of size l as a subgraph, then π contains the pattern σ by construction. This is because if we consider the permutation matrix representation of π and delete all columns except the ones that correspond to separating and encoding entries for vertices of $K_l \subseteq G$, we get a matrix that differs from the permutation matrix of σ only by the additional presence of columns that encode the connection of the vertices of K_l to the vertices outside of K_l. By deleting these columns (and the empty rows resulting from the above deletions) we arrive at the permutation matrix representation of σ, implying σ is a pattern of π.

For the other direction, assume there is a function $\phi : [|\sigma|] \to [|\pi|]$ certifying the pattern. We start by noting that there are no decreasing subsequences of length $\frac{1}{4}z$ in π avoiding all separating runs. This is because such a sequence contains at most one entry from $C(v)$ for each $v \in V(G)$. At the same time, it cannot simultaneously contain an entry from $C(u)$ and $C(v)$ for u, v chosen from different connected components. This is because the construction of the encoding permutation places vertices from the same component consecutively and the entries encoding a component placed earlier in the ordering have strictly smaller values than those from a later component. This bounds the length of the subsequence by $n' < \frac{1}{4}z$.

Furthermore, *any* decreasing subsequence of π contains entries from at most one pair of separating runs. This is because once the sequence hits a separating run of a vertex v, all its subsequent entries can only be from the pair of separating runs of v. Any decreasing subsequence therefore starts with less than $\frac{1}{4}z$ encoding entries, which are then followed by entries of a pair of separating runs of some vertex.

We now show that the certifying function ϕ naturally leads to a mapping from $V(K_l)$ to $V(G)$. Consider any vertex $v \in V(K_l)$. The function ϕ maps the subsequence of σ formed by the pair of separating runs of v to a decreasing subsequence of π of the same length. As argued, such a long decreasing subsequence starts with less than $\frac{1}{4}z$ encoding entries of π, which are then followed by at least $\frac{7}{4}z$ entries from a pair of separating runs of some vertex $u \in V(G)$. This implies that the middle entry $L(v)$ of the left separating run of v in σ needs to be mapped by ϕ to the left separating run of $u \in G$. Additionally, the middle entry $R(v)$ of the right separating run of v in σ needs to be mapped by ϕ somewhere in the right separating run of the same vertex u. The above establishes a mapping from $V(K_l)$ to $V(G)$ denoted by f_ϕ.

We claim f_ϕ to be a graph homomorphism. Fix any pair of vertices v_1, v_2 of K_l such that $v_1 < v_2$. We show that $f_\phi(v_1), f_\phi(v_2)$ are connected by an edge in G. Since there is an edge between v_1 and v_2 in K_l, Observation 8 implies that the set of values $\sigma[L(v_1), R(v_1)]$ contains precisely one number p with $\sigma(R(v_2)) \leq p \leq \sigma(L(v_2))$. Since ϕ certifies the pattern σ in π, there needs to be an entry of π with an index between $\phi(L(v_1))$ and $\phi(R(v_1))$ and value between $\pi(\phi(R(v_2)))$ and $\pi(\phi(L(v_2)))$. Observation 8 then implies there is an edge between $f_\phi(v_1)$ and $f_\phi(v_2)$. Thus, f_ϕ is a homomorphism and G contains a clique of size l. □

The above reduction can be directly used within the cross-composition framework to show our result:

Proof of Theorem 6. We set L to be the set of all pairs (K_l, G), where K_l is a clique, G is a connected graph containing K_l as a subgraph. It is widely known that deciding $x \in L$ is NP-complete.

We introduce a cross-composition of L into Permutation Pattern Matching. Let R be an equivalence relation on $\{0,1\}^*$ with the following properties: the binary sequences that are not representing a pair (K_l, G), where K_l is a clique and G a graph, are placed in a single equivalence class designated for malformed input sequences; a pair of strings representing instances (K^1, G^1) and (K^2, G^2), respectively, is related in R if and only if $|V(K^1)| = |V(K^2)|, |V(G^1)| = |V(G^2)|$. Clearly, R is a polynomial equivalence relation. For instances $(K_l, G_1), (K_l, G_2), (K_l, G_3), \ldots, (K_l, G_t)$ from the same equivalence class of R, we produce an instance of the Permutation Pattern Matching problem where we ask if $\pi_z(K_l)$ is in $\pi_z(G)$, where G is a disjoint union of graphs $G_1, \ldots, G_t$ and z is set to $4 \cdot |V(G_1)| + 4$. Lemma 9 shows that the answer to this problem is YES if and only if at least one of the instances (K_l, G_i) belongs to L. Since the parameter of the pattern matching instance is $|\pi_z(K_l)|$, which can be bounded by $|V(G_i)| \cdot 2z + |V(G_i)|^2$ for any i, we can apply Theorem 5. □

4 Conclusion

Guillemot and Marx [9] have shown that the Permutation Pattern Matching problem can be solved in $2^{\mathcal{O}(\ell^2 \log \ell)} \cdot n$ time. They raised the question of whether a faster FPT algorithm could be obtained and outlined a strategy for achieving this using their notion of decompositions of permutations. This relied on the bound from the Stanley-Wilf conjecture not being tight. Fox [8] has shown this is not the case. (Still, [8] gives an improved $2^{\mathcal{O}(\ell^2)} \cdot n$ algorithm.) The non-existence of a polynomial kernel is a further indication of the difficulty of the problem.

References

[1] S. Ahal, Y. Rabinovich: On Complexity of the Subpattern Problem, SIAM J. Discrete Math., 22(2), 629-649, 2008.

[2] H. L. Bodlaender, R. G. Downey, M. R. Fellows, D. Hermelin: On problems without polynomial kernels, J. Comput. Syst. Sci., 75(8), 423-434, 2009.

[3] H. L. Bodlaender, B. M. P. Jansen, S. Kratsch: Cross-Composition: A New Technique for Kernelization Lower Bounds, Proc. 28th STACS, 165–176, 2011.

[4] M.-L. Bruner, M. Lackner: The computational landscape of permutation patterns, Pure Mathematics and Applications, 24 (2), 83-101, 2013.

[5] M.-L. Bruner, M. Lackner: A fast algorithm for permutation pattern matching based on alternating runs, F. V. Fomin and P. Kaski, editors, SWAT, 7357, Lecture Notes in Computer Science, 261-270, Springer, 2012.

[6] P. Bose, J. F. Buss, A. Lubiw: Pattern Matching for Permutations, Inf. Process. Lett., 65(5), 277-283, 1998.

[7] L. Fortnow, R. Santhanam: Infeasibility of instance compression and succinct PCPs for NP, J. Comput. Syst. Sci., 77(1), 91-106, 2011.

[8] J. Fox: Stanley-Wilf limits are typically exponential, arXiv:1310.8378, 2014, to appear in Advances in Mathematics.

[9] S. Guillemot, D. Marx: Finding Small Patterns in Permutations in Linear Time, Proc. SODA 2014, 2014.

[10] L. Ibarra: Finding pattern matchings for permutations, Information Processing Letters, 61(6), 293–295, 1997.

[11] M. Klazar: The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture, Formal Power Series and Algebraic Combinatorics, 250-255, 2000.

[12] D. E. Knuth: The Art of Computer Programming, Vol. 1: Fundamental Algorithms, 2nd edition, Addison-Wesley, 1973.

[13] L. Lovász: Combinatorial Problems and Exercises, North-Holland, 1979.

[14] A. Marcus, G. Tardos: Excluded permutation matrices and the Stanley-Wilf conjecture, J. Comb. Theory, Ser. A, 107(1), 153-160, 2004.

[15] D. Rotem: Stack-sortable permutations, Discrete Math, 33, 185-196, 1981.

[16] C. Schensted: Longest increasing and decreasing subsequences, Classic Papers in Combinatorics, 299–311, 1987.

[17] R. Simion, F. W. Schmidt: Restricted permutations, European J. Combinatorics, 6, 383-405, 1985.

[18] H. Wilf's address to the SIAM meeting on Discrete Mathematics in 1992.